Scalable Anomaly Detection in Batch Polishing Processes for Inertial Confinement Fusion Shells

Shashank Galla [*1] **Akash Tiwari** [*1] **Kshitij Bhardwaj** [*2] **Sean Michael Hayes** [2] **Satish T.S Bukkapatnam** [1] **Suhas Bhandarkar** [2]

Abstract

In the domain of Inertial Confinement Fusion (ICF), ensuring the surface quality of target shells is critically demanding, as process anomalies can lead to significant losses of both materials and time. This paper presents a scalable anomaly detection methodology tailored specifically for the specialized task of batch polishing these shells. Our approach utilizes an autoencoder model trained on normal process data, applying an Exponentially Weighted Moving Average (EWMA) control chart to the model's reconstruction loss to detect anomalies. This methodology effectively handles large datasets and meets the rigorous surface quality standards required to maximize ignition yield in ICF operations. We validate our methodology on datasets from two distinct batch polishing experiments, H200 and H202, which encompass multiple hours of operation. The paper also details the tuning of hyperparameters for the EWMA control chart and includes both an ablation study and a comparative analysis with other anomaly detection methods. Our results demonstrate that this methodology outperforms existing approaches by promptly detecting subtle anomalies with minimal delay and low rates of false positives. Specifically, it achieves an average detection delay of 1.3 seconds for the H200 dataset and 19.5 seconds for the H202 dataset, thereby contributing significantly to the advancement of efficient and effective anomaly detection in the ICF domain.

[1]Department of Industrial and Systems Engineering, Texas A&M University, College Station, TX 77843. [2]Lawrence Livermore National Lab, Livermore, California 94550.. Correspondence to: Shashank Galla <shashankgalla@tamu.edu>, Kshitij Bhardwaj <bhardwaj2@llnl.gov>.

Proceedings of the AI for Science workshop at 41^{st} International Conference on Machine Learning, Vienna, Austria (2024). Copyright 2024 by the author(s).

1. Introduction

Decades of sustained research in laser-mediated nuclear fusion have recently culminated in demonstration of ignition, where the output energy is greater than the laser input energy. At the National Ignition Facility (NIF) in December 2022, an experiment yielded about 1.5 times greater fusion energy than optical energy, which has generated much interest in developing nuclear fusion as a highly concentrated energy source that produces no harmful by-products (Abu-Shawareb et al., 2024). With multiple more demonstrations that followed, it is now hoped that nuclear fusion could one day revolutionize the field of clean energy and greatly alleviate the global energy crisis (Nature, 2024).

The approach used at NIF for this demonstration is called Inertial Confinement Fusion (ICF). Here, a roughly 2mm spherical shell made of high-density carbon (HDC), which is akin to diamond (Biener et al., 2009), houses the deuterium-tritium fuel on the interior required for ICF. The fusion process is initiated by 192 laser beams, which impart energy to the shell. This causes the HDC shell to implode rapidly towards the fuel core, compressing and heating it to the levels (1e9 bars and 1e8 K) needed to overcome the Coulombic repulsion between the deuterium and tritium nuclei.

Ensuring the optimal use of deposited laser energy heavily depends on the surface quality of the shells (Zylstra et al., 2022). Any defects in the surface of the shell can lead to significant instabilities, disrupting the uniform and symmetric implosion necessary for successful fusion (Hurricane et al., 2023; Casey et al., 2015; Schmitt et al., 2013). The stringent requirements for the surface quality of the shells necessitate ultra-high precision finishing processes, which is comprised of a multi-stage polishing process. After undergoing fabrication and polishing, only those shells that meet the high surface quality standards are eligible for a fusion shot.

Polishing processes for batches of approximately 20 shells are inherently time-consuming and susceptible to significant anomalies, which can result in catastrophic events (Lawrence Livermore National Laboratory, 2024). These anomalies typically originate from unanticipated deviations within the polishing system or from surface irregularities such as crack initiation sites on the shells. Such cracks

may propagate, causing fractures and spalling of the HDC coating. The disintegration of even a single shell can lead to debris that compromises the integrity of the entire batch, resulting in considerable losses.

To mitigate these catastrophic events and safeguard the polishing process, implementing robust online process monitoring is crucial. By detecting incipient anomalies, this system allows operators to halt the process timely, thereby preventing extensive damage to expensive batches of shells (Bukkapatnam et al., 2019). Advanced side-channel sensors, including vibration sensors, acoustic emission sensors, and microphones, are employed due to their high sensitivity to subtle sub-micrometer level surface changes (Jin et al., 2022).

However, the challenge lies in managing the vast volumes of data generated by these sensors during the extended, multi-hour polishing sessions. The sensors when sampled at high frequencies, accumulate substantial datasets that demand efficient real-time processing and data management solutions. Addressing this issue, our work aims to develop a scalable anomaly detection methodology capable of handling such large-scale datasets effectively.

In this work, we explore the use of an accelerometer to gather extensive vibration data from the polishing process, coupled with scalable machine learning techniques for early anomaly detection. Specifically, an autoencoder model reconstructs segments of the vibration signal, capitalizing on its ability to learn normal data patterns and detect deviations (Sakurada & Yairi, 2014). A high-pass filter is applied to the reconstruction loss to improve the detection capability and reduce cyclicity and noise. The processed signal is then monitored using an Exponentially Weighted Moving Average (EWMA) control chart, an effective method for identifying shifts in process behavior (Hunter, 1986; Montgomery, 2007).

Our methodology is implemented on two experimental datasets, H200 and H202, each spanning multiple hours. The system demonstrates a rapid detection capacity, with lags of 1.3 seconds for H200 and 19.5 seconds for H202, facilitating timely interventions to prevent potential damage. This work significantly advances the application of artificial intelligence in ICF and precision manufacturing. By integrating autoencoders with EWMA control charts, we enhance process efficiency and prevent catastrophic failures, broadening the spectrum of AI applications within the nuclear and precision manufacturing sectors.

The paper is structured as follows: Section 2 delves into the technical background of autoencoder models and EWMA control charts. Section 3 describes the methodology for detecting anomalies in the batch shell polishing process. Section 4 discusses the tuning of hyperparameters for the EWMA control chart and presents findings from a case study involving two batch shell polishing experiments, including an ablation study and benchmarking comparisons. Section 5 concludes the paper and outlines avenues for future research.

2. Background

In this study, we introduce a dual-stage framework for anomaly detection tailored for complex manufacturing processes. The initial stage focuses on identifying critical features for monitoring using a machine learning model. We select an autoencoder for this purpose due to its demonstrated effectiveness in unsupervised anomaly detection. In the subsequent stage, these identified features are monitored using an Exponentially Weighted Moving Average (EWMA) control chart to detect out-of-control (OOC) events indicative of anomalies. This paper further elaborates on the specific autoencoder architectures employed for unsupervised feature extraction and the application of EWMA control charts for the detection of anomalies.

2.1. Autoencoder

Autoencoders have demonstrated efficacy in various anomaly detection applications (Yan et al., 2023). These neural network models operate on the principle of distinguishing between data with unseen properties and data with seen properties during training. Autoencoders consist of two main components: an encoder (f_{θ_e}) that compresses the input $\mathbf{w}_k \in \mathbb{R}^n$ into a latent space representation $\mathbf{z} \in \mathbb{R}^m$, and a decoder ($g_{\theta_d}$) that reconstructs the input from this latent representation to produce $\hat{\mathbf{w}}_k$. Mathematically, the encoder-decoder structure can be expressed as follows:

$$\mathbf{z} = f_{\theta_e}(\mathbf{w}_k) \tag{1}$$

$$\hat{\mathbf{w}}_k = g_{\theta_d}(\mathbf{z}) \tag{2}$$

An input is classified as anomalous if the reconstruction loss $\mathcal{L}_k$, typically measured by the Mean Squared Error (MSE), exceeds a predefined threshold. The loss function is given by:

$$\mathcal{L}(w_k, \hat{w}_k) = \frac{1}{n}\sum_{i=1}^{n}(w_{k,i} - \hat{w}_{k,i})^2 \tag{3}$$

In our approach, rather than using a predefined threshold, we implement an EWMA control chart, as detailed in Sec. 2.2, to detect anomalies. This method enhances sensitivity to subtle shifts in the monitored feature. For our specific anomaly detection problem, the model is trained exclusively on vibration data from normal polishing runs. This allows the model to learn the typical characteristics of normal operations, facilitating the discrimination between normal and

anomalous conditions in the polishing process. Also, a high-pass filter is applied to $\mathcal{L}_k$ to eliminate cyclicity and noise from the data.

2.2. EWMA control charts

Control charts serve as a fundamental instrument in statistical process control, designed to establish upper and lower control limits and plot data points over time. These charts are pivotal for monitoring processes and identifying out-of-control (OOC) conditions when these limits are breached (Montgomery, 2009). Statistical process control is typically segmented into two phases: Phase I involves setting the initial control limits, while Phase II uses these established limits to monitor ongoing data.

Among various control chart formats, the Exponentially Weighted Moving Average (EWMA) control chart stands out for its efficacy in handling autocorrelated time series data (Shi et al., 2022). In this study, we utilize the EWMA control chart to monitor the filtered reconstruction loss ($\mathcal{L}'_k$) from an autoencoder model. This choice is informed by the autocorrelation present in the time series data under analysis.

The EWMA method offers a nuanced mechanism for process monitoring by assigning incrementally greater weight to more recent observations, thereby enhancing the detection of subtle changes and anomalies in the process. This methodological approach ensures that recent data points exert a more significant influence on the control measures, facilitating timely identification of deviations and maintaining the integrity of process monitoring.

The Phase I process involves obtaining data points ($\mathcal{L}'_k$), setting parameters (α, L, W), computing statistics (z_t, $\overline{X}$ and $\hat{\sigma}$), and establishing control limits (LCL/UCL). If data points fall outside the control limits, they are removed, and control limits are recalculated. Once stable control limits are established, they are used for monitoring in Phase II. The EWMA Statistic (z_t) is given by:

$$\mathrm{z}_t = \alpha\mathcal{L}'_t + (1-\alpha)\mathrm{z}_{t-1}, \forall t = 1, \ldots, T \tag{4}$$

where z_0 is data-point at time t = 0, z_t is data-point at time t and α is the smoothing parameter which corresponds to the weight of historic observations. The steady-state control limits are given by

$$LCL/UCL = \overline{X} \pm L\hat{\sigma}\sqrt{\frac{\alpha}{2-\alpha}} \tag{5}$$

where $\overline{X}$ is the average of the phase I observed data, $\hat{\sigma}$ is the standard deviation of the phase I observed data and L is the control limit co-efficient. α and L are the two parameters that determine the design of an EWMA chart. In our methodology, Phase I involves analyzing a steady-state process of length W, where W serves as a critical design parameter. The rationale and selection criteria for W are discussed in detail in Section 4.2.

3. Methodology

This section outlines our proposed methodology, which is structured into three key phases: Data collection, Feature extraction, and Process monitoring. A detailed schematic of our methodology for anomaly detection in batch polishing processes is depicted in Figure 1.

3.1. Data collection

3.1.1. Experimentation

Polishing experiments are performed on a batch of 20 shells, polished simultaneously. Vibration data, denoted as y_t for $t \in \{1, \ldots, T\}$, is captured from these experiments using an accelerometer (Kistler 8728A500), where T indicates the total duration of the experiment. The data collection occurs at a sampling rate of 10 kHz, equivalent to every 0.0001 seconds.

3.1.2. Data labeling

The experimental data y_t is segmented using a window size β. Each window, represented as $w_k = \{y_t, \ldots, y_{t+\beta-1}\}$, undergoes auditory analysis by a human operator who listens to the sounds produced by the vibration signals during the polishing process. The segments are labeled as $g_k \in \{0, 1\}$, with 0 indicating normal operating conditions and 1 signifying anomalous conditions. These labels are assigned based on auditory characteristics and changes in signal amplitudes, allowing for the detection of anomalies that manifest as unusual sound patterns or significant amplitude variations. This method enhances the detection of subtle deviations that might not be evident in purely numerical analysis.

3.2. Feature extraction

3.2.1. Autoencoder

An autoencoder model is trained on windowed data (w_k, g_k) where $g_k = 0$ represents normal operating conditions. During training, the model learns to compress the input data and subsequently reconstruct it, optimizing the parameters to minimize the discrepancy between the input and output. After training, the model generates reconstructed data $\hat{w}_k$ for each window w_k. The reconstruction loss, denoted as $\mathcal{L}(w_k, \hat{w}_k)$, quantifies the differences between the original and reconstructed data. This metric is pivotal for anomaly detection, as higher loss values suggest potential deviations

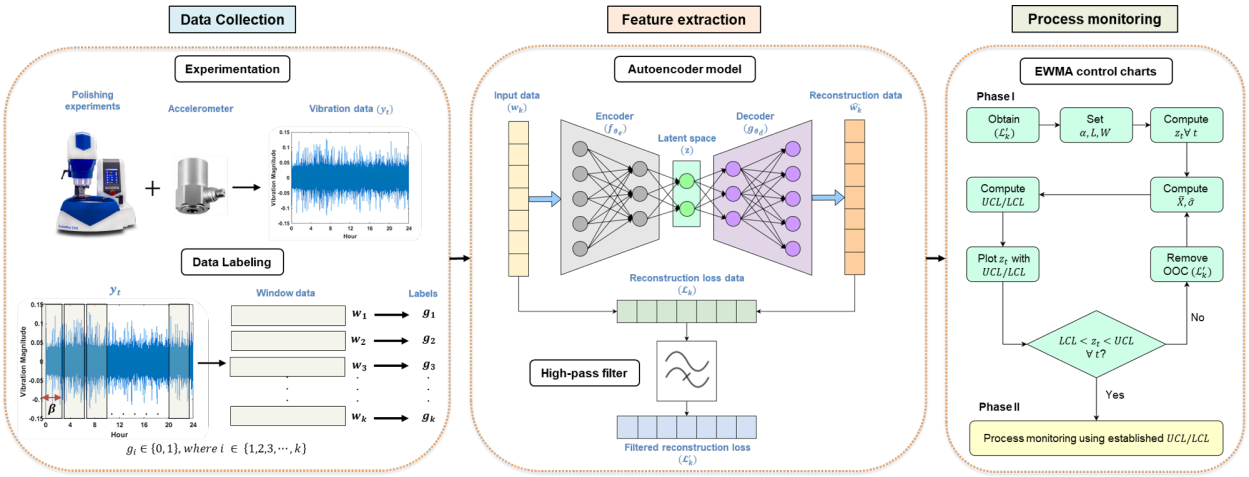

Figure 1. Proposed Methodology for detecting anomaly during batch polishing process

from normal process conditions.

3.2.2. High-pass filter

Following the model training, the reconstruction loss data $\mathcal{L}_k$ is subjected to high-pass filtering, a crucial step to eliminate cyclical patterns and reduce noise, enhancing the model's sensitivity to subtle anomalies. This filtering effectively attenuates low-frequency fluctuations, which might otherwise mask the true dynamics of the process. By emphasizing high-frequency components, this approach ensures that the data processed in subsequent analysis stages more accurately represents significant deviations.

3.3. Process monitoring

In Phase I of the process monitoring, an Exponentially Weighted Moving Average (EWMA) control chart is developed, establishing both lower (*LCL*) and upper (*UCL*) control limits, using the high-pass filtered reconstruction loss $\mathcal{L}'_k$ from the autoencoder model. Phase II involves the application of this EWMA control chart to detect out-of-control data points indicative of anomalies. Any data point that breaches these established control limits is classified as an anomaly.

4. Batch polishing case study

In general, process changes can be classified into two categories: major and subtle changes. Major changes, or process anomalies, significantly alter the dynamics of the system. The primary objective of this case study is to identify the first major process anomaly, which marks a critical transition in operational behavior. This initial anomaly is particularly crucial as it heralds a substantial shift in the process dynamics, often manifested by an increased amplitude in the signal post-anomaly. Our proposed methodology is specifically designed to quickly detect this first major change, facilitating timely intervention to prevent and mitigate potential damage.

4.1. Dataset description

As previously discussed, we conduct polishing experiments and collect extensive vibration datasets, with two such instances designated as *H200* and *H202*. These experiments last $T_{H200} = 2.78$ hours and $T_{H202} = 6.94$ hours, respectively. Each experiment culminates in a catastrophic event that results in the complete decimation of the shells, leading to substantial damage to the entire batch.

Figure 2 displays the time portraits of the vibration signals from these experiments, highlighting the occurrence of the first process anomaly. The vibration data $y_t \mid t \in \{1, \ldots, T\}$ are segmented into windows of size $\beta = 1000$ samples, which corresponds to a time resolution of 0.1 seconds. A human operator examines each window w_k by listening to the playback of the signals at a sampling rate of 10 kHz. This auditory review facilitates the labeling of time portrait regions as normal ($g_k = 0$) or anomalous ($g_k = 1$), as discussed in Section 3.1.2.

The labeled data from the aforementioned experiments is utilized for training the autoencoder.

4.2. Autoencoder training and filtering

For training our autoencoder, only a subset of the *H200* dataset is utilized, specifically excluding the initial 100 seconds identified as sensor startup noise by auditory analysis. Subsequently, 900 seconds of data, categorized as normal by the human operator, are used for training. The remaining data from *H200* are allocated for validation and testing pur-

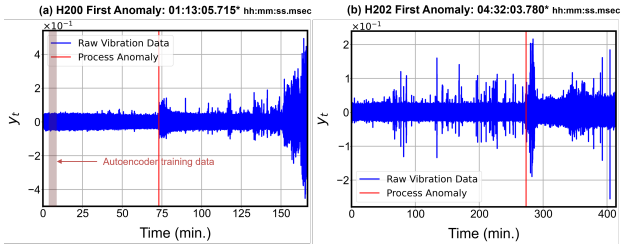

Figure 2. (a) *H200* vibration time portrait with training data for autoencoder (red patch) (b) *H202* vibration time portrait

poses. The trained model is also assessed using the *H202* dataset.

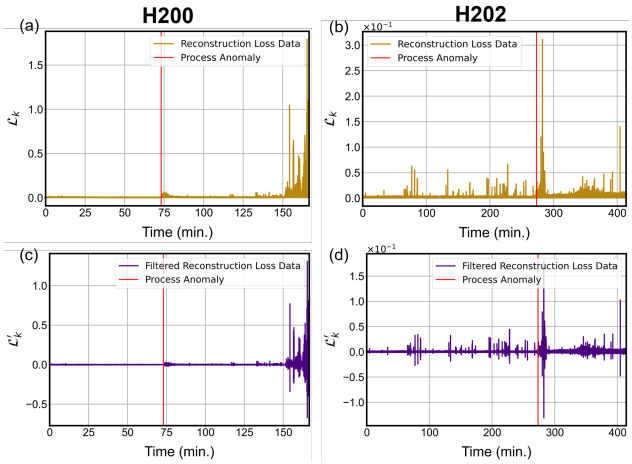

Figure 3. (a) Reconstruction loss from autoencoder for *H200* (b) Reconstruction loss from autoencoder for *H202*. (c) Filtered reconstruction loss from autoencoder for *H200* (cutoff frequency: 1000 Hz) and (d) for *H202* (cutoff frequency: 1000 Hz).

The architecture of our autoencoder is specifically designed to discern complex patterns within the datasets. It incorporates fully connected layers, enhanced with batch normalization and ReLU activation functions to ensure training stability and facilitate convergence. The encoder compresses the input data from 1000 dimensions to a 100-dimensional latent space, and the decoder restores the data back to its original dimensionality. This dimensionality reduction is crucial as it ensures that only the most relevant features are preserved. Figure 3(a,b) displays the autoencoder's reconstruction loss for the *H200* and *H202* test datasets, with the loss distribution highlighting significant anomalous regions—two in *H200* and one in *H202*—where the losses significantly exceed the baseline.

After completing the training phase, the reconstruction loss $\mathcal{L}_k$ is extracted from the autoencoder. To further process this data, a high-pass filter with a cutoff frequency of 1000 Hz is applied. The rationale behind choosing this particular cutoff frequency is elaborated in Section 4.5. Figures 3(c,d) depict the filtered reconstruction loss from the autoencoder for the *H200* and *H202* datasets, respectively.

Table 1. Training Configuration Parameters

Parameter	**Value**
Loss Function	MSE
Optimizer	Adam
Learning Rate	0.001
Beta Parameters	(0.99, 0.999)
Weight Decay	10^{-5}
Training Epochs	500

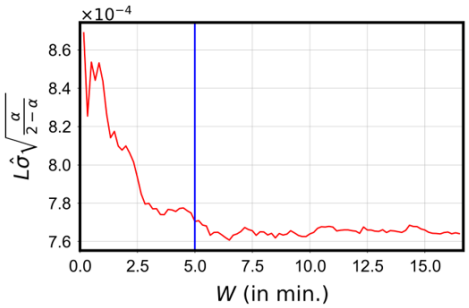

Figure 4. Choice of Phase I window size W is chosen based on convergence of control limit width. We set W = 5 min after which the control limit width is asymptotic.

4.3. Design parameter settings for EWMA control chart

In the subsequent stage, after deriving $\mathcal{L}'_k$, the analysis transitions to Phase I. This initial phase is critical for selecting the appropriate parameters: the window size W, the smoothing constant α, and the control limit L. It is a common practice to set α within the range [0.1, 0.3]. Opting for a lower α value, such as $\alpha = 0.1$, enhances sensitivity to subtle changes, thereby facilitating the detection of significant shifts in the process behavior. We select $L = 3$ to balance the reduction of false positives with the retention of sensitivity to significant variations, accommodating the high variability observed in $\mathcal{L}'_k$.

The determination of W during Phase I is established through a numerical experiment. The control bounds $\frac{UCL-LCL}{2}$ stabilize at $W = 5$ minutes, as evidenced by the variations shown in Fig. 4. Consequently, W is set to 5 minutes for Phase I analysis to effectively detect process anomalies.

4.4. Results and discussions

Following the completion of hyperparameter analysis and the establishment of fixed parameters, Phase I analysis is

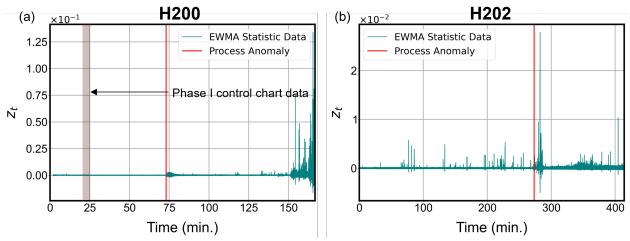

Figure 5. (a) EWMA statistic time portraits for *H200* corresponding window size used in phase I for control chart development are illustrated as brown patch (b) EWMA statistic time portraits for *H202*

Table 2. Timing and adjustment of the first anomalies detected in the datasets, accounting for system lag.

Dataset	First Anomaly (hh:mm:ss)	Detection (hh:mm:ss)	Lag (sec.)
H200	01:13:05	01:13:07	1.3
H202	04:32:03	04:32:23	19.5

commenced. We construct Exponentially Weighted Moving Average (EWMA) control charts for the *H200*'s $\mathcal{L}'_k$. To ensure data stationarity, the initial 20 minutes are excluded, and control limits (Upper Control Limit, UCL, and Lower Control Limit, LCL) are derived from the subsequent 5 minutes of data. These established limits are subsequently utilized to detect anomalies in the unseen datasets of *H200* and *H202* during Phase II. Figure 5(a,b) displays the EWMA statistic constructed on $\mathcal{L}'_k$ for *H200* and *H202*, respectively.

The timing of the first detected anomalies is documented in *Table 2*. An anomaly is recognized when a data point surpasses the established control limits, indicating an out-of-control (OOC) condition. This breach constitutes the detection event, and the time interval from the actual occurrence of the anomaly to its detection is termed the detection lag.

For the datasets *H200* and *H202*, the recorded detection lags were 1.3 seconds and 19.5 seconds, respectively. These metrics underscore the efficiency of the EWMA charts in swiftly detecting anomalies, though responsiveness varies across the datasets. The performance for *H202* may be further improved through domain adaptation strategies, such as those proposed by (Ganin & Lempitsky, 2015), to accommodate the differing conditions encountered in polishing operations between H200 and H202.

Figure 6 illustrates the EWMA control charts applied to *H200*. This implementation resulted in two false positives, depicted in Figure 6(b). The sensitivity of EWMA charts to minor variations can be modulated by adjusting the smoothing constant α. Setting α to 0.1, as previously discussed, resulted in these false positives. Adjusting α directly affects

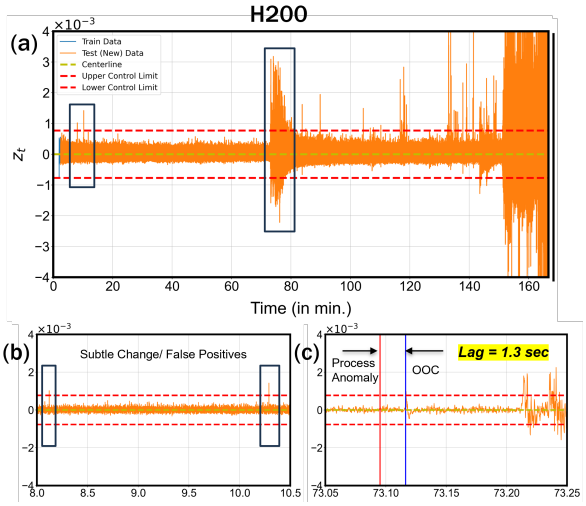

Figure 6. (a) EWMA control chart of $\mathcal{L}_k$ Phase II Anomaly Detection for *H200*. (b) illustrate subtle change/false positives and (c) illustrate the lag between the ground truth occurrence of anomaly and time when control chart detects anomaly.

the detection dynamics: increasing α heightens sensitivity, potentially increasing false positives, while decreasing α tends to delay the detection of anomalies. Figure 6(c) further displays the original process anomalies alongside points identified as out of control (OOC), underscoring the capability of EWMA charts to discern process anomalies effectively.

4.5. Ablation study

To evaluate the impact of different components on the performance of our anomaly detection system, we conducted an ablation study utilizing four distinct analytical frameworks. This study was implemented using the H200 dataset. The configurations along with their respective detection latencies are summarized in Table 3.

- **Impact of autoencoder:** Employing an autoencoder in conjunction with a high-pass filter (cutoff frequency: 1000 Hz) facilitated anomaly detection with a minimal lag of 1.3 seconds. In contrast, omitting the autoencoder increased the detection lag to 6 seconds for H200 and 2 seconds for H202, albeit with no false positives recorded at either cutoff frequency. Despite the absence of false positives, our primary concern remains with minimizing detection lag; the occurrence of two false positives is deemed acceptable within this context. This finding illustrates the critical role of the autoencoder as a non-linear filter in enhancing the sensitivity of reconstruction loss data to process anomalies.
- **Impact of cutoff frequency:** Increasing the cutoff

Table 3. Detailed results showing the detection lag and false positives for various analysis frameworks.

Sl. No	Analysis Framework	Lag (sec.)	False positives count
1	Vibration + Autoencoder + High pass Filter (1000 Hz) + EWMA	(+)1.3	2
2	Vibration + Autoencoder + High pass Filter (2500 Hz) + EWMA	(+)7.2	2
3	Vibration + High pass Filter (1000 Hz) + EWMA	(+)7.3	0
4	Vibration + High pass Filter (2500 Hz) + EWMA	(+)9.7	0

Table 4. Performance comparison of different methods based on lag and false positives.

Sl. No	Method	Lag (sec.)	False positives count
1	Autoencoders + EWMA	**1.3**	**2**
2	IForest	52	11
3	LOF	94	9
4	KNN	52	10
5	HBOS	50	7

frequency from 1000 Hz to 2500 Hz induced an additional lag of 2.5 seconds. This increase likely attenuates essential signal components integral to anomaly detection, underscoring the significance of selecting an appropriate cutoff frequency for optimal performance of the anomaly detection system.

4.6. Comparative analysis

To evaluate the effectiveness of our proposed methodology, we conducted a comparative analysis using established methods tailored to our problem domain. The *H200* dataset served as the basis for this comparative analysis. Four prevalent methods in anomaly detection research were selected for their proven efficacy: Isolation Forest (IForest) (Liu et al., 2008), Local Outlier Factor (LOF) (Breunig et al., 2000), k-Nearest Neighbors (KNN) (Ramaswamy et al., 2000), and Histogram-based Outlier Score (HBOS) (Goldstein & Dengel, 2012). The results are presented in Table 4. Evaluation metrics included detection lag and the number of false positives/subtle changes, with the objective of minimizing both. As shown in Table 4, our proposed method surpasses all considered benchmarks in these metrics.

5. Conclusions and future work

This work introduces a scalable anomaly detection methodology that leverages an EWMA control chart for autoencoder reconstruction loss. Designed to accommodate large, extensive datasets such as *H200* and *H202*, which encompass multiple hours of data, our approach is particularly effective in managing the complexities of batch polishing spherical shells. The robustness of the polishing process is enhanced by demonstrating the methodology's capability in handling large-scale data environments.

Our approach is engineered to detect process anomalies at the earliest possible instance, facilitating timely interventions. We achieved detection latencies of 1.3 seconds for *H200* and 19.5 seconds for *H202*, indicating that such delays do not compromise the integrity of production batches. Instead, the swift identification and remediation of anomalies contribute to maintaining operational efficiency and minimizing potential disruptions.

The proposed method exhibits heightened sensitivity to subtle changes, occasionally leading to false positives as depicted in Figure 6(b). This sensitivity arises from the model's acute responsiveness to minor signal fluctuations. While this attribute is beneficial for early anomaly detection, it also heightens the risk of mistaking benign variations for potential faults.

The timely detection of anomalies is critical in conserving valuable resources. A notable challenge in batch shell polishing is the precise temporal localization of concurrent damage across multiple shells, which complicates the interpretation of vibration signals. In response, our trained autoencoder model serves as a tool for pinpointing anomalies within vibration data, thereby facilitating further investigative efforts into the nature of these anomalies.

Future work will incorporate the use of cameras to precisely locate temporal anomalies, reducing reliance on human-generated labels. Further optimization of the autoencoder and EWMA control chart parameters (α, L) is necessary to enhance the system's sensitivity to subtle process deviations, often overlooked by human operators, and to improve detection lag and enable proactive maintenance strategies. Additionally, domain adaptation techniques will be investigated to refine the prediction accuracy across different polishing settings. We also aim to implement our anomaly detection technique on edge devices located adjacent to the sensors for in-situ monitoring during the early stages of anomaly formation. This deployment will facilitate immediate notifications before potential disruption to the entire batch of shells, thus preserving the integrity of the manufacturing process.

Broader impact

The methodology developed herein offers a modular, general-purpose approach for time-series based anomaly

detection in process monitoring. It enables the integration of various machine learning models and control charts to adjust detection sensitivity according to specific needs. Applied to batch polishing, this technique allows for the timely detection of catastrophic damages, potentially saving significant production hours. Additionally, an online monitoring system can proactively alert operators to halt machinery and address detected faults, thereby enhancing operational efficiency and reducing costly downtime. This approach underscores its potential to significantly improve safety and productivity in manufacturing environments.

Acknowledgements

This work was performed under the auspices of the U.S. Department of Energy by LLNL under contract DE-AC52 -07NA27344 and was supported by the LLNL laboratory-directed research and development (LDRD) program under project 23-ERD-014.

This material is partially based on work supported while serving at the National Science Foundation. Any opinion, findings, conclusions or recommendation expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

During the preparation of this work, the author(s) used ChatGPT, developed by Open AI, to improve readability and overall language. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content. No AI tool was used for problem formulation or analysis or to draw insights in the research process.

References

Abu-Shawareb, H., Acree, R., Adams, P., Adams, J., Addis, B., Aden, R., Adrian, P., Afeyan, B., Aggleton, M., Aghaian, L., et al. Achievement of target gain larger than unity in an inertial fusion experiment. *Physical Review Letters*, 132(6):065102, 2024.

Biener, J., Ho, D., Wild, C., Woerner, E., Biener, M., El-Dasher, B., Hicks, D., Eggert, J., Celliers, P., Collins, G., et al. Diamond spheres for inertial confinement fusion. *Nuclear Fusion*, 49(11):112001, 2009.

Breunig, M. M., Kriegel, H.-P., Ng, R. T., and Sander, J. Lof: identifying density-based local outliers. In *Proceedings of the 2000 ACM SIGMOD international conference on Management of data*, pp. 93–104, 2000.

Bukkapatnam, S. T., Afrin, K., Dave, D., and Kumara, S. R. Machine learning and ai for long-term fault prognosis in complex manufacturing systems. *CIRP Annals*, 68(1): 459–462, 2019.

Casey, D., Milovich, J., Smalyuk, V., Clark, D., Robey, H., Pak, A., MacPhee, A., Baker, K., Weber, C., Ma, T., et al. Improved performance of high areal density indirect drive implosions at the national ignition facility using a four-shock adiabat shaped drive. *Physical review letters*, 115 (10):105001, 2015.

Ganin, Y. and Lempitsky, V. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pp. 1180–1189. PMLR, 2015.

Goldstein, M. and Dengel, A. Histogram-based outlier score (hbos): A fast unsupervised anomaly detection algorithm. *KI-2012: poster and demo track*, 1:59–63, 2012.

Hunter, J. S. The exponentially weighted moving average. *Journal of quality technology*, 18(4):203–210, 1986.

Hurricane, O., Patel, P., Betti, R., Froula, D., Regan, S., Slutz, S., Gomez, M., and Sweeney, M. Physics principles of inertial confinement fusion and us program overview. *Reviews of Modern Physics*, 95(2):025005, 2023.

Jin, S., Tuo, R., Tiwari, A., Bukkapatnam, S., Aracne-Ruddle, C., Lighty, A., Hamza, H., and Ding, Y. Hypothesis tests with functional data for surface quality change detection in surface finishing processes. *IISE transactions*, pp. 1–17, 2022.

Lawrence Livermore National Laboratory. Target evolution: Key to llnl's continued success, 2024. URL https://www.llnl.gov/article/49721/target-evolution-key-llnls-continued-success.

Liu, F. T., Ting, K. M., and Zhou, Z.-H. Isolation forest. In *2008 eighth ieee international conference on data mining*, pp. 413–422. IEEE, 2008.

Montgomery, D. C. *Introduction to statistical quality control*. John Wiley & Sons, 2007.

Montgomery, D. C. *Statistical quality control*, volume 7. Wiley New York, 2009.

Nature. Nuclear fusion breakthrough: this physicist helped to achieve the first-ever energy gain. *Nature*, 625:11–12, 2024. doi: 10.1038/d41586-023-04045-8. URL https://www.nature.com/articles/d41586-023-04045-8.

Ramaswamy, S., Rastogi, R., and Shim, K. Efficient algorithms for mining outliers from large data sets. In *Proceedings of the 2000 ACM SIGMOD international conference on Management of data*, pp. 427–438, 2000.

Sakurada, M. and Yairi, T. Anomaly detection using autoencoders with nonlinear dimensionality reduction. In *Proceedings of the MLSDA 2014 2nd workshop on machine learning for sensory data analysis*, pp. 4–11, 2014.

Schmitt, M. J., Bradley, P. A., Cobble, J. A., Hsu, S. C., Krasheninnikova, N. S., Kyrala, G. A., Magelssen, G. R., Murphy, T. J., Obrey, K. A., Tregillis, I. L., et al. Defect-induced mix experiment for nif. In *EPJ Web of Conferences*, volume 59, pp. 04005. EDP Sciences, 2013.

Shi, Z., Mamun, A. A., Kan, C., Tian, W., and Liu, C. An lstm-autoencoder based online side channel monitoring approach for cyber-physical attack detection in additive manufacturing. *Journal of Intelligent Manufacturing*, pp. 1–17, 2022.

Yan, S., Shao, H., Xiao, Y., Liu, B., and Wan, J. Hybrid robust convolutional autoencoder for unsupervised anomaly detection of machine tools under noises. *Robotics and Computer-Integrated Manufacturing*, 79:102441, 2023.

Zylstra, A., Kritcher, A., Hurricane, O., Callahan, D., Ralph, J., Casey, D., Pak, A., Landen, O., Bachmann, B., Baker, K., et al. Experimental achievement and signatures of ignition at the national ignition facility. *Physical Review E*, 106(2):025202, 2022.